Recognising the Last Record of a Sequence

Alexander Gnedin[*]

February 8, 2008

Abstract

We study the best-choice problem for processes which generalise the process of records from Poisson-paced i.i.d. observations. Under the assumption that the observer knows distribution of the process and the horizon, we determine the optimal stopping policy and for a parametric family of problems also derive an explicit formula for the maximum probability of recognising the last record.

1 Introduction

Maximising the probability of stopping at the extreme of a sequence of random marks is the classical objective in sequential decision problems widely known as the best-choice or 'secretary' problems [3, 13, 22]. Problems of this kind can be formulated in terms of the embedded process of records, because the overall extreme (e.g. minimum) is the *last record* observation.

In a basic version of the problem introduced by Gilbert and Mosteller [10, Section 3] the marks are sampled at discrete times from the uniform distribution, and the objective of the observer is to stop at the minimum among the first n marks. The sequence of values of sequential minima, called *lower records*, undergoes a *stick breaking* process $X_1, X_1 X_2, X_1 X_2 X_3, \ldots$, where the X_j's are independent copies of a prototypical random factor X whose distribution is uniform. Given the record values, the *durations* of records are independent, and for r a generic record value, the duration of a record with this value has geometric distribution with parameter r. See [1, 16, 19] for these basic facts of the theory of random records. The optimal policy in the stopping problem of [10] is rather complicated, as it involves a sequence of thresholds for which no closed-form expression is available, and for the same reason there is no explicit formula for the optimal probability.

According to another version of the problem, the marks are observed at epochs of a unit Poisson process, and the goal is to stop at the minimum mark before given horizon T (see [12] and references therein). This problem allows much more explicit results: the optimal policy prescribes stopping at the first time the record process breaks through a hyperbolic boundary, and there is an explicit formula for the optimal probability. The continuous time problem corresponds to the model sometimes called *Poisson-paced records*

---

[*]Insitute of Mathematics, Utrecht University, The Netherlands, gnedin@math.uu.nl

[16, Section 9], the difference with the discrete-time model is that the duration of a record with value r has exponential distribution with rate r. For large n and T the discrete and continuous time versions are close to the same limiting form introduced in [11], in particular the limiting optimal probability is given by the formula first obtained by Samuels [21] in discrete setting. These and related results are reviewed in Section 7.3.

In this paper we consider the continuous-time problem of recognising the last record under a more general assumption that the occurences of records follow a stick-breaking scheme, with factor X having an arbitrary distribution on the unit interval. Models of this kind appear in many contexts such as branching processes, search problems, sequential packing problems and random partitions [2, 4, 8, 20]. Although we just postulate the behaviour of records without any reference to some more rich observable process, the model in focus is related to one concept of sequential extreme for sampling from certain partially ordered spaces, including spaces $\mathbb{R}^d$ with continuous product distributions. This connection is detailed in Section 3.

We will show that the optimal policy is always of the same form as in the case of uniform X. In one special case of parametric family of beta distributions we express the optimal probability in terms of the incomplete gamma function. In general, however, it does not seem possible to write a closed-form expression for the stopping value. Still, we argue that under minor side conditions on the law of X, as $T \to \infty$, there exists a limiting value which may be interpreted as the optimal probability of recognising the last record in a stopping problem with infinitely many observations. The famous best-choice probability benchmark e^{-1} will show up as a sharp lower bound.

2 The model

We shall model the occurences of records by means of a nonincreasing right-continuous Markov process $R = (R_t, t \geq 0)$ with the following type of behaviour: given the current state is $r > 0$, the process jumps at rate r to a new state rX, where X is a prototypical random factor with a given distribution in the open interval $]0, 1[$. In the event t is a jump instant of R we say that a record occurs at time t and intepret R_t as the *weight* of the record. The weights of consequitive records decrease, while the sojourns of R, which include the first record time and further durations of records, are stochastically increasing.

In more detail, the weights of records undergo stick-breaking $r_0X_1, r_0X_1X_2, \ldots$, where X_j's are independent replicas of X and $r_0 = R_0$ is the initial state of R. The sequence of sojourns may be represented as

$$E_1/r_0, E_2/(r_0X_1), E_3/(r_0X_1X_2), \ldots$$

where E_j's are i.i.d. unit exponential variables, independent of the X_j's. Thus, conditionally given the weights of records, the sojourns are independent exponential variables.

We are interested in the problem of maximising the probability of recognising the last record of R before a given horizon T, by means of a nonanticipating policy (stopping time) adapted to the natural right-continuous filtration of R. For π such a policy the efficiency is measured by the probability that π is a record time not exceeding T and that

no further record occurs before time T:

$$\mathbb{P}(R_{\pi-} > R_\pi = R_T, \pi < T) = \mathbb{E}\left[\exp\left\{-(T-\pi)R_\pi\right\}\mathbf{1}(R_{\pi-} > R_\pi, \pi < T)\right], \qquad (1)$$

where the second expression involves the adapted probability of recognising the last record when the stopping occurs.

In the terminology going back to Gilbert and Mosteller [10], this stopping problem should be qualified as a problem with 'full information', meaning that the observer learns the weights of records and knows their distribution exactly. Under 'no-information' problem we understand the optimal stopping problem where only policies based on record times are allowed.

3   Chain records

Sampling from arbitrary continuous distribution F leads to the stick-breaking process for records with uniform X. This is seen by defining the weight via $v \mapsto F(v)$ and by noting that this mapping preserves the ranking and transforms a sample from F into a sequence of uniform variables. In this section we discuss some extensions of this framework.

Sampling from certain discrete distributions also leads to stick-breaking process for records. Define a distribution by allocating the geometric masses pq^{k-1} (where $p + q = 1$ and $0 < p < 1$) at points of some decreasing sequence z_k, $k = 1, 2, \ldots$. Consider *strict* lower records in a sample from such distribution. Define the weights by means of the *left*-continuous distribution function $v \mapsto F(v-)$. If the first sample value is z_k, then the next observation is a record with probability q^k; from this we see that the weights of records follow the stick-breaking scheme with factor

$$X =_d \sum_{k=1}^{\infty} pq^{k-1}\delta_{q^k},$$

where δ_x is the Dirac mass at x and $=_d$ denotes the equality in distribution.

Sampling from other distributions on reals is not consistent with the stick-breaking model for records. We will look now in higher dimensions.

Consider $\mathbb{R}^d$ endowed with some continuous product distribution μ and the natural strict partial order $\prec$. For a sample $V_1, V_2, \ldots$ from $(\mathbb{R}^d, \mu)$, we say that a *chain record* occurs at index j if either $j = 1$ or $j > 1$ and V_j is $\prec$-smaller than the last chain record in the sequence $V_1, \ldots, V_{j-1}$. Define the *weight* of a chain record by means of the multivariate distribution function $v \mapsto \mu\{u \in \mathbb{R}^d : u \prec v\}$. The weights of chain records follow a stick-breaking process with the density $\mathbb{P}(X \in \mathrm{d}x)/\mathrm{d}x = |\log x|^{d-1}/(d-1)!$ for the factor X. Indeed, the componentwise probability transform establishes isomorphism between the ordered probability space $(\mathbb{R}^d, \mu, \prec)$ and the unit cube $[0, 1]^d$ with the Lebesgue measure, which implies that the law of X is the same as the distribution of the product of d independent uniform variables, whence the formula for the density.

Chain records in $\mathbb{R}^d$ were introduced in [14]. Unlike other kinds of multidimensional records surveyed in [17], the chain records cannot be regarded as 'generalised minima', because permutations of $V_1, \ldots, V_{j-1}$ may destroy or create a chain record at index j.

3

The sequence of chain-record marks is a 'greedy' decreasing chain in the partially ordered sequence of marks, in the sense that element V_j is joined to the chain each time when the monotonicity constraint is not violated (as to be compared, e.g., with the longest chain among the first n marks).

The definition of chain record extends in an obvious way to sampling from an arbitrary Borel space $\mathcal{Z}$ endowed with a probability measure μ and a measurable strict partial order $\prec$. The weights are defined by means of the function $v \mapsto \mu(L_v)$, where $L_v = \{u \in \mathcal{Z} : u \prec v\}$ is the lower section of $\prec$ at $v \in \mathcal{Z}$. Call the space $(\mathcal{Z}, \mu, \prec)$ *lower-homogeneous* if (i) $\mu(L_v) > 0$ for μ-almost all points $v \in \mathcal{Z}$, and (ii) the lower section L_v with conditional measure $\mu(\cdot)/\mu(L_v)$ is isomorphic, as a partially ordered probability space, to the whole space $(\mathcal{Z}, \mu, \prec)$. Since all L_v's are in this sense the same, the weights of chain records in a sample from a lower-homogeneous space undergo a stick-breaking with the factor $X =_d \mu\{u \in \mathcal{Z} : u \prec V\}$ where V has distribution μ.

It is easily seen that $[0,1]^d$ with uniform distribution is a lower-homogeneous space, hence this is true also for $\mathbb{R}^d$ with continuous product distribution. Another example is the *interval space* which has intervals $]a, b[\subset [0,1]$ as elements, the partial ordering $\prec$ defined by inclusion, and a measure $\mu(\mathrm{d}a\mathrm{d}b) = \alpha(\alpha - 1)(b - a)^\alpha \mathrm{d}a\,\mathrm{d}b$ (with parameter $\alpha > 1$); in this case $\mathbb{P}(X \in \mathrm{d}x)/\mathrm{d}x = (\alpha - 1)(x^{-1/\alpha} - 1)$. Although both examples are instances of Bollobás-Brightwell *box-spaces* [6] (which have all intervals $\{u : v \prec u \prec w\}$ for $v \prec w$ isomorphic to the whole space and not only L_v's), there are many other lower-homogeneous spaces that are not box-spaces. By the transitivity of partial order, the distribution of X appearing in this way must satisfy the inequality $\mathbb{P}(X \le x) \ge x$, $x \in [0,1]$.

4   Stopping the embedded Markov chain

A fundamental property of the process R is *self-similarity*: for each $r > 0$, the law of $(R_t,\ t \ge 0)$ with $R_0 = r$ is identical to the law of $(rR_{rt},\ t \ge 0)$ given $R_0 = 1$. This implies that, when the law of X is fixed, the 'size' of the problem is determined by a single parameter $r_0 T$.

Self-similarity is a clue to derive the optimal policy. If stopping has not occur before and including time t and if the current state is $R_t = r$, then the conditional optimal stopping problem is equivalent to the unconditional problem with initial state 1 and horizon $(T - t)r$. This motivates associating with R (with fixed parameters r_0, T and the law for X) another decreasing Markov process $B = (B_t,\ t \ge 0)$,

$$B_t = (T - t)_+ R_t, \quad t \ge 0,$$

with the initial state $B_0 = r_0 T$ and the absorbing terminal state 0. Obviously, it is sufficient to consider the policies adapted to B, with understanding that the last record before T corresponds to the last jump of B before absorption at 0. The sequence of locations visited by B at the record times is a discrete-time homogeneous Markov chain which follows the transition scheme $s \mapsto (s - E)_+ X$, $s \ge 0$, where E is a rate-1 exponential variable independent of X.

In terms of the embedded chain the optimal policy is determined in a standard way, by comparing two kinds of risk. If the current state of B is s, the probability that no

further records occur is $p_0(s) = e^{-s}$. On the other hand, the probability that exactly one record will occur is

$$p_1(s) = \int_0^s e^{-t}\, \mathbb{E}\left[p_0((s-t)X)\right] \mathrm{d}t = e^{-s}\, \mathbb{E}\left[\frac{e^{s(1-X)} - 1}{1 - X}\right].$$

Inspecting two extremes $s = 0$ and ∞ and exploiting monotonicity, we see that the equation

$$\mathbb{E}\left[\frac{e^{s(1-X)} - 1}{1 - X}\right] = 1 \tag{2}$$

has a unique positive solution s_*. Because

$$p_0(s) < p_1(s) \iff s > s_*,$$

and because B has decreasing paths (until getting absorbed) we are in the familiar *monotone case* of optimal stopping, hence the optimal policy stops at the first jump of B within the region $[0, s_*]$. Translating this back in terms of R we see that it is optimal to stop at the first record time when the condition $(T - t)R_t \le s_*$ is satisfied. In particular, if $Tr_0 \le s_*$ it is optimal to stop at the very first record.

More generally, for $s > 0$ we denote π_s the policy which prescribes stopping at the first record time when $(T - t)R_t \le s$ holds. Summarising the above discussion we conclude:

Proposition 1. *The policy π_{s_*} with s_* satisfying* (2) *is optimal.*

Assuming $R_0 = 1$, let $v(T, s)$ be the value of the policy π_s, i.e. the probability that exactly one record before T satisfies $(T - t)r \le s$. (By self-similarity the case of arbitrary $R_0 = r_0$ can be reduced to that.) Obviously,

$$v(T, s) = p_1(T), \quad \text{for } T < s. \tag{3}$$

The first-record decomposition readily yields an integral equation

$$v(T, s) = \int_0^T \mathbb{E}\left[v(X(T - t), s)\,\mathbf{1}((T - t)X > s) + e^{-(T-t)X}\,\mathbf{1}((T - t)X \le s)\right] e^{-t} \mathrm{d}t,$$

which for $s = s_*$ is the familiar dynamic programming equation for the optimal value. In the differential form this becomes

$$\partial_T\, v(T, s) = -v(T, s) + \mathbb{E}\left[v(TX, s)\,\mathbf{1}(TX > s)\right] + \mathbb{E}\left[e^{-TX}\,\mathbf{1}(TX \le s)\right]. \tag{4}$$

The equation (4) is of delayed type, which only in exceptional cases admits a closed-form solution. For instance, when the distribution of X is δ_x, the solution is a piecewise-analytical function which should be computed recursively in the intervals $T \in [s/x^{k-1}, s/x^k]$ for $k = 1, 2, \ldots$, starting from $[0, s/x]$ where $v(T, s) = p_1(s)$ holds.

The collection of sites which B visits at record times is not a Poisson process, since otherwise $v(T, s)$ were constant in T for $T > s$. It is therefore surprising that the maximum of $v(T, s)$ in s is attained at the same point s_*, for all $T > s_*$.

5 The lower bound

Suppose for a while that the law of X is δ_1. In this case (4) is easily solved as $v(T, s) = (T \wedge s) e^{-(T \wedge s)}$. Thus $s_* = 1$ and for $T \geq 1$ the optimal probability is $v(T, s_*) = e^{-1}$, which also coincides with the maximum of $p_1(s) = se^{-s}$. To bring this conclusion into the familiar 'no-information' framework note that $R_t \equiv 1$, hence there is no updating of record weights. For the same reason, the record times are the epochs of a unit Poisson process, hence the stopping problem amounts to recognising the last Poisson epoch on $[0, T]$, which is the 'no-information' problem for Poisson process due to Browne [7]. A characteristic feature of this case is that $v(T, s)$ is constant in T for $T > s$ (see the last remark in Section 4).

We show next that the familiar benchmark $e^{-1} = 0.367\ldots$ yields a universal lower bound in our model.

Proposition 2. *For every distribution of X the optimal probability satisfies $v(T, s_*) > e^{-s_*}$ for $T > s^*$. Above that, $s_* < 1$ hence*

$$v(T, s_*) > e^{-1} \quad \text{for} \ \ T > 1,$$

and this bound is sharp.

Proof. Suppose $r_0 T > s_*$. The process B can enter $[s_*, 0]$ by either continuously drifting down or jumping down through s_*. In the first case the conditional probability of sucess with π_* is $p_1(s_*) = p_0(s_*) = e^{-s_*}$. In the second case this probability is $\mathbb{E}[e^{-S}] > e^{-s_*}$, with some random $S < s_*$, because π_{s_*} will stop. The estimate readily follows.

Applying the inequality $e^{1-x} > 1 + (1 - x)$ for $0 < x < 1$, we see that the left-hand side of (2) is larger than 1 for $s = 1$, therefore the root satisfies $s_* < 1$. The bound e^{-1} is approached by letting the law of X to approach δ_1. $\qquad\square$

The same argument yields a more general inequality $v(T, s) > \min(p_0(s), p_1(s))$ for $T > s$, where the right side assumes the largest value at $s = s_*$.

For X uniform $s_* = 0.804\ldots$ and the lower bound is $e^{-s_*} = 0.447\ldots$, while for X with density $|\log x|$ these are $0.743\ldots$ and $0.475\ldots$.

6 Entrance from the infinity

For asymptotic considerations we shall vary the initial state and denote $\mathbb{P}_r$ the law of R with $R_0 = r$. Assume that $\mathbb{E} |\log X| < \infty$ and that the distribution of X is not supported by a geometric progression (note that these are precisely the conditions for applicability of the renewal theorem [9] to $-\log X$). Let

$$f(\lambda) = \mathbb{E}[X^\lambda]$$

be the Mellin transform of X. Clearly, $-f'(0) = \mathbb{E} |\log X|$. Adapting [5, Theorem 1] we have:

Proposition 3. *Under the above assumptions, as $r \to \infty$, the law $\mathbb{P}_r$ has a weak limit $\mathbb{P}_\infty$ characterised by $R_t =_d Y/t$, $t > 0$, where Y is a random variable uniquely determined by its moments*

$$\mathbb{E}\left[Y^k\right] = \frac{1}{-f'(0)} \prod_{j=1}^{k-1} \frac{j}{1-f(j)}, \quad k = 1, 2, \ldots \tag{5}$$

Corollary 4. *Under these circumstances there exists a limit $v(\infty, s_*) = \lim_{T \to \infty} v(T, s_*)$ which is the maximum probability of recognising the last record for the process $(R_t, t \leq 1)$ under $\mathbb{P}_\infty$.*

Proof. This follows from the form of the optimal policy and the fact that the point process of sites visited by B at record times has a weak limit as $B_0 \to \infty$. $\square$

The law of Y determined by (5) can be considered as a kind of extreme-value distribution. For instance, Y is exponential for X uniform, while Y is distributed like the product of independent uniform and exponential variables for X with density $|\log x|$.

Denoting τ_1, τ_2, w_1, w_2 the times and weights of the last record and the record before the last, the performance of π_s in the infinite problem can be written as

$$v(\infty, s) = \mathbb{P}_\infty((1-\tau_1)\rho_1 < s < (1-\tau_2)\rho_2) = \mathbb{P}_\infty((1-\tau_1)\rho_1 < s) - \mathbb{P}_\infty((1-\tau_2)\rho_2 < s). \tag{6}$$

In principle, the moments (5) determine the distribution of these variables, for instance

$$\mathbb{P}_\infty(\tau_1 < t) = \mathbb{E}\left[e^{Y(1-1/t)}\right],$$

but it seems impossible to use this for writing $v(\infty, s)$ in some explicit form.

7   The beta case

We proceed with more concrete computations under the assumption that the distribution of X is beta$(\theta, 1)$, with the density

$$\mathbb{P}(X \in \mathrm{d}x)/\mathrm{d}x = \theta x^{\theta-1}, \quad x \in [0,1],$$

where θ is a positive parameter. The instance $\theta = 1$ corresponds to the uniform distribution. This class of stick-breaking processes has a feature that under $\mathbb{P}_\infty$ both the range of R and the point process of record times are Poisson point processes with intensity measure $\theta\,\mathrm{d}z/z$, $z > 0$. The law of R_1 under $\mathbb{P}_\infty$ is a gamma distribution.

The integral

$$p_1(s) = \int_0^1 \frac{e^{-sx} - e^{-s}}{1-x}\,\theta x^{\theta-1}\,\mathrm{d}x$$

does not simplify, hence it should be included in the final formula for the optimal probability as it is.

7.1 Computing the value

For $T > s$ a substitution translates (4) into

$$T^\theta \, \partial_T \, v(T, s) = -T^\theta v(T, s) + \int_s^T v(t, s)\theta t^{\theta-1}\mathrm{d}t + \int_0^s e^{-t}\theta t^{\theta-1}\mathrm{d}t\,.$$

Differentiating in T and simplifying we are lead to

$$Tg'' + (T + \theta)g' = 0 \qquad\qquad (7)$$

for $g(T) = v(T, s)$. Solving this and taking into account the boundary condition at $T = s$ yields

$$v(T, s) = \Gamma(-\theta + 1, s, T)e^s s^\theta p_1'(s) + p_1(s)\,, \quad \text{for } T > s\,, \qquad\qquad (8)$$

where

$$\Gamma(a, b, c) = \int_b^c e^{-t}t^{a-1}\mathrm{d}t$$

denotes the incomplete gamma function, and

$$p_1'(s) = -p_1(s) + \frac{\theta}{s^\theta}\Gamma(\theta, 0, s).$$

For the optimal s_* using $p_1(s_*) = e^{-s_*}$ we obtain from (8)

$$v(T, s_*) = \Gamma(-\theta + 1, s_*, T)[-s_*^\theta + e^{s_*}\theta\Gamma(\theta, 0, s_*)] + e^{-s_*}\,, \qquad\qquad (9)$$

which is the optimal probability of stopping at the last record. The formula is valid for $T \geq s^*$. The optimal probability $v(\infty, s_*)$ in the limit problem is just obtained taking $T = \infty$ in the integral in (9), which reads as a generalised exponential integral function

$$\Gamma(-\theta + 1, s_*, \infty) = \int_{s_*}^\infty \frac{e^{-t}}{t^\theta}\,\mathrm{d}t\,.$$

The following table shows some numerical values of this probability computed with a help of `Mathematica`.

θ	0.1	0.25	0.5	1	2	5	20
s_*	0.709	0.731	0.760	0.804	0.857	0.922	0.976
$v(\infty, s_*)$	0.913	0.814	0.703	0.580	0.481	0.410	0.377

The data suggest to examine the extreme values of the parameter θ.

As $\theta \to \infty$ the beta distribution approaches δ_1, hence $s_* \to 1$ and $v(\infty, s_*) \to e^{-1}$. Thus the beta family may be interpreted as a bridge between the 'full-information' problem ($\theta = 1$) and the 'no-information' problem ($\theta = \infty$).

Note that, for arbitrary $\theta > 0$, in consequence of the Poisson character of the record times under $\mathbb{P}_\infty$, the time-threshold policy $\pi = \min\{t > T/e, R_t > R_{t-}\}$ yields the limit probability of success equal e^{-1} for $T \to \infty$.

As $\theta \to 0$ the beta distribution approaches δ_0. In this regime the optimal s_* approaches $\log 2$. Selecting T_0 sufficiently large to secure occurence of at least one record with probability at least $1 - \epsilon$, and then sending θ to 0, we will have $v(T_0, s_*) \geq 1 - \epsilon$, because with high probability exactly one record occurs before horizon T. For $T > T_0$ (9) implies $v(T, s_*) > v(T_0, s_*)$, therefore the trivial upper bound $v(\infty, s_*) < 1$ is sharp as the law of X varies.

8

7.2 A smooth fit

With the explicit formula (8) in hand we can alternatively characterise s_* as the maximiser of $v(T, s)$ in s. Equating $\partial_s v(T, s)$ to 0 we see then that s_* is a root of the equation

$$sp_1''(s) + (s + \theta)p_1'(s) = 0, \tag{10}$$

which is, in fact, equivalent to $p_0(s) = p_1(s)$ due to the identity

$$sp_1''(s) + (s + \theta)p'(s) = \theta(p_0(s) - p_1(s)).$$

Comparing (10) with (7) shows that two branches of $v(T, s_*)$, for $T \leq s_*$ and $T \geq s_*$, match at $T = s_*$ together with *two* derivatives. This degree of smoothness is characteristic for $s = s_*$, as is also seen by the following argument. Write the probability of success with policy π_s as

$$v(T, s) = \int_0^s p_1'(t)dt + \int_s^T \partial_T v(t, s)dt,$$

and note that the optimisation of s amounts to finding a 'switch' which maximises the sum of integrals. Inspecting the monotonicity properties of the integrands

$$p_1'(t) = \frac{e^{-t}}{t^\theta}(p_1'(t)t^\theta e^t) \quad \text{and} \quad \partial_T v(t, s) = \frac{e^{-t}}{t^\theta}(p_1'(s)s^\theta e^s)$$

shows that the maximum is achieved if they are tangential at the switching location, which is precisely the condition (10). Thus s_* is indeed the only value of s such that $\partial_T v(T, s)$ has no break at $T = s$.

7.3 The uniform case

For completeness we bring together known formulas for the case of uniform factor X.

The solution to

$$\int_0^s \frac{e^t - 1}{t}dt = 1$$

has the approximate value $s_* = 0.804\ldots$ The limit probability

$$v(\infty, s_*) = (e^{s_*} - s_* - 1)\int_{s_*}^\infty \frac{e^{-t}}{t}\,dt + e^{-s_*} = 0.580\ldots,$$

was obtained first numerically in [10] by interpolation from discrete-time problems, derived in [21] from (6), and shown [3] by some series computations with the Poisson process. The analogous formula for $v(s, T)$ with finite T appeared in [12].

The process of records under $\mathbb{P}_\infty$ corresponds to the set of $\prec$-minimal atoms of a unit-rate Poisson point process on $\mathbb{R}_+^2$ (recall that an atom is $\prec$-minimal if there are no other Poisson atoms south-west of it). The unique properties of the planar Poisson process allow more delicate computations. Under $\mathbb{P}_\infty$ the density of π_s is [12]

$$\mathbb{P}(\pi_s \in dt)/dt = \frac{t-1}{t}(e^{-ts} - e^{-ts/(1-t)}) + s\Gamma\left(0, st, \frac{st}{1-t}\right) + 1 - e^{-st}, \quad t \in [0, 1],$$

9

This integrates to some number less than 1 because with positive probability π_s does not stop at all (this probability is aproximately $0.1995\ldots$ for the optimal policy π_{s_*}). The optimal probability can be also represented as the integral

$$v(\infty, s_*) = \int_0^1 w(t)\mathrm{d}t\,,$$

with $w(t)$ the *winning rate*, equal to the chance that π_{s_*} stops correctly in time $\mathrm{d}t$. The graph of w was sketched long ago [10, Figure 3], and the following explicit formula for the winning rate is a recent result [15]:

$$w(t) = -e^{-s_*} + \frac{e^{-s_* t} - e^{-s_* t/(1-t)}}{t} + \frac{e^{-s_* t} - t e^{-s_*}}{1-t} +$$
$$\frac{s_*}{1-t}\left[\Gamma\left(0, s_*, \frac{s_*}{1-t}\right) - \Gamma\left(0, s_* t, \frac{s_* t}{1-t}\right)\right],$$

(the boundary values $w(0) = 1 - e^{-s_*}$, $w(1) = e^{-s_*}$ were indicated in [10]).

8   Concluding remarks

A discrete-time version of the problem with fixed horizon n is associated with a process analogous to R but with geometric durations of records. The optimal policy is known only for uniform X. Moreover, it is not clear if the monotone case of optimal stopping applies for the general distribution of X. Using techniques from [18] one can show that under the assumptions of Section 6 the discrete-time problem can be aproximated, for n large, by the limiting problem with continuous time, hence the policy 'stop at index j if a record occurs with weight r satisfying $r(n - j) \leq s_*$' is asymptotically optimal.

It would be also interesting to evaluate suboptimal policies like 'stop at the first record with weight below given w' or 'stop at the first record that occurs after a given time t_0'. This is not so easy in general since such policies are not adapted to B.

References

[1] Arnold, B.C., Balakrishnan, N. and Nagaraja, H.N. (1998) *Records*, NY, Wiley.

[2] Baryshikov, Yu.M. and Gnedin, A.V. (2001) Counting intervals in the packing process, *Ann. Appl. Probab.* 11: 863–877.

[3] Berezovsky, B.A. and Gnedin, A.V. (1984) *The best choice problem*, Moscow, Nauka.

[4] Bertoin, J. (2006) *Random fragmentation and coagulation processes*, forthcoming book.

[5] Bertoin, J. and Caballero, M.-E. (2002) Entrance from 0+ for increasing semi-stable Markov processes, *Bernoulli* 8: 195-205.

[6] Bollobás, B. and Brightwell, G. (1991) Box spaces and random partial orders, *Trans. Amer. Math. Soc.* 124: 59-72.

[7] Browne, S. (1994) Records, mixed Poisson processes and optimal selection: an intensity approach, Columbia Univ. working paper.

[8] Devroye, L. (1999) Universal limit laws for depths in random trees, *SIAM J. Comp.* 28: 409-432.

[9] Feller, W. (1971) *An Introduction to Probability Theory and its Applications*, vol. 2, NY, Wiley.

[10] Gilbert, J. and Mosteller, F. (1966) Recognizing the maximum of a sequence, *J. Amer. Stat. Assoc.* 61: 35-73.

[11] Gnedin, A.V. (1996) On the full-information best-choice problem, *J. Appl. Prob.* 33: 678-687.

[12] Gnedin, A.V. (2004) Best choice from the planar Poisson process, *Stoch. Proc. Appl.* 111: 317-354.

[13] Gnedin, A.V. (2005) Objectives in the best-choice problems, *Sequential Analysis* 24: 1-11.

[14] Gnedin, A. (2005) Counting the chain records, available at arXiv.

[15] Gnedin, A. and Miretskiy, D. (2005) Winning rate in the full information best-choice problem, available at arXiv.

[16] Goldie, C.M. and Bunge, J. (1999) Record sequences and their applications, in *Handbook of Statistics vol. 19*, (*Stochastic Processes: Theory and Methods*, Shanbhag, D.N. and Rao, C.R. eds), pp. 277-308, North Holland, Amsterdam.

[17] Goldie, C.M. and Resnick, S. (1989) Records in a partially ordered set, *Ann. Probab.* 17: 678-699.

[18] Kühne, R. and Rüschendorf, L. (2000) Approximation of optimal stopping problems, *Stoch. Proc. Appl.* 90: 301-325.

[19] Nevzorov, V.B. (2001) *Records*, Transl. Math. Monographs, Providence, AMS.

[20] J. Pitman, Exchangeable and partially exchangeable random partitions (1995) *Prob. Th. Rel. Fields* 102: 145-158.

[21] Samuels, S.M. (1982) Exact solutions for the full information best choice problem, *Purdue Univ. Stat. Dept. Mimeo Series 82-17.*

[22] Samuels, S.M. (1991) Secretary problems. Chapter 16 of *Handbook of Sequential Analysis* (B.K.Ghosh and P.K. Sen eds), NY, Marcel Dekker.